


Shedding Light on Distributed System Executions

<http://bestchai.bitbucket.org/shiviz/>

Inspecting distributed system logs is tedious and error-prone:

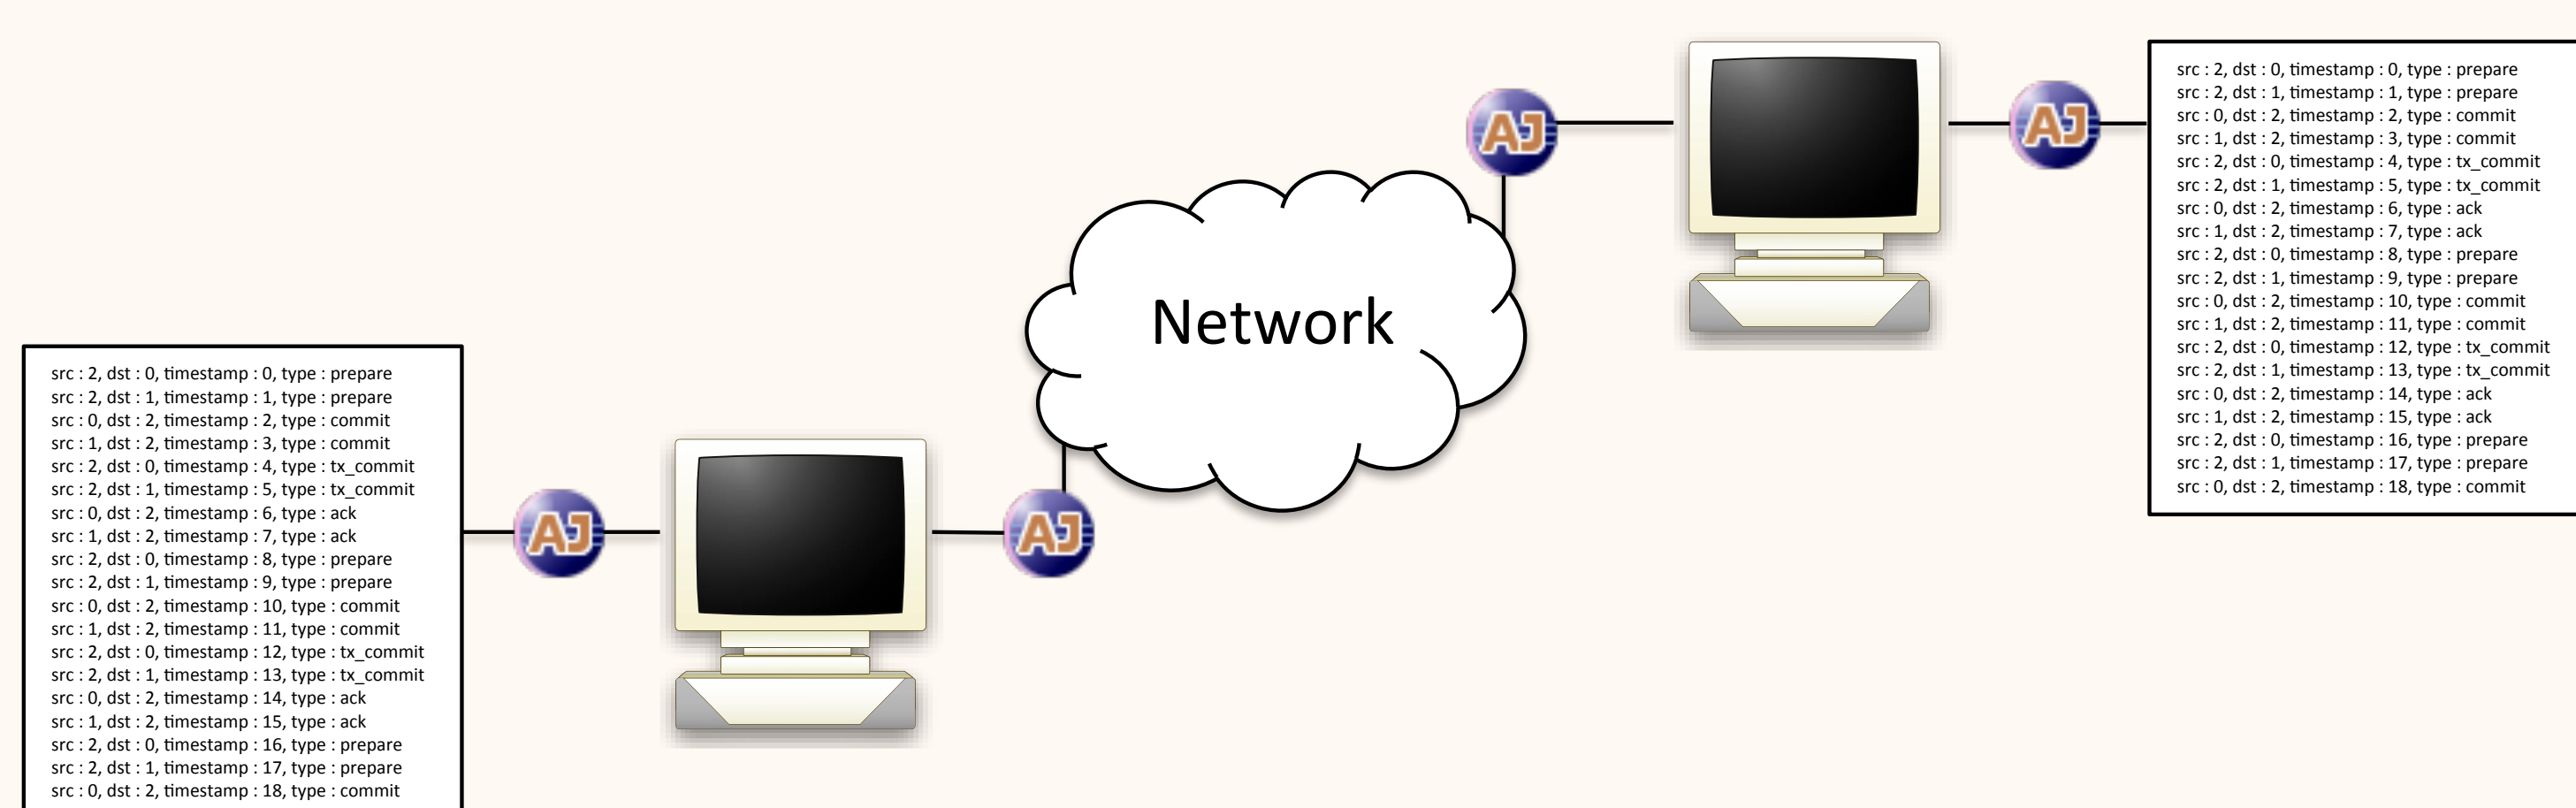


Understand and debugging requires knowing how nodes communicated:

- Were events **X** and **Y** concurrent, or did one precede another?
- Did node **A** ever communicate with node **B**? When?
- Did node **B** ever communicate with node **C**? When?
- Can we effectively visualize series of communication?

ShiVector tool

- Automatically instruments a distributed system logging with ordering information (vector clocks)
- Transparently interposes on all socket communication and logging (via AspectJ)



System

System + ShiVector

```
...
INFO Timeline request
GET /timeline
INFO New status
POST status="Lunch"
...
```

```
...
INFO Timeline request [ 3, 10, 7, 2 ]
GET /timeline [ 3, 11, 7, 2 ]
INFO New status [ 5, 12, 10, 2 ]
POST status="Lunch" [ 5, 13, 10, 2 ]
...
```

Ongoing work

- Visualizing multiple executions
- Mining behavior motifs from large logs
- Query/filter language for needle-in-a-haystack analysis
- Scaling visualization to long executions/many nodes

ShiViz tool

- Concisely visualizes a ShiVector log of a distributed execution
- Lightweight (runs in a browser)
- Reveals the communication topology in the system (gives developers a new perspective on familiar console logs)

Execution log

```
[2013-05-24 23:28:00,637 voldemort.store.metadata.MetadataStore] INFO
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:2)
[2013-05-24 23:28:00,749 voldemort.store.metadata.MetadataStore] INFO
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:3)
[2013-05-24 23:28:00,830 voldemort.server.VoldemortServer] INFO Using NIO Connector.
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:4)
[2013-05-24 23:28:00,845 voldemort.server.VoldemortServer] INFO Using NIO Connector for Admin Service.
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:5)
[2013-05-24 23:28:00,845 voldemort.server.VoldemortService] INFO Starting voldemort-server
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:6)
[2013-05-24 23:28:00,903 voldemort.server.VoldemortServer] INFO Starting 7 services.
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:7)
[2013-05-24 23:28:00,903 voldemort.server.VoldemortService] INFO Starting storage-service
42795@jvoldemortThread[main,5,main] ("42795@jvoldemortThread[main,5,main]:8)
Examples: Facebook data-center log Voldemort log SimpleDB log Facebook log with multiple executions
```

ShiVector log input

Visualization

